

Representing Data by a Mixture of Activated Simplices

Chunyu Wang
Peking University

John Flynn
UCLA

Yizhou Wang
Peking University

Alan L. Yuille
UCLA

Abstract

We present a new model which represents data as a mixture of simplices. Simplices are geometric structures that generalize triangles. We give a simple geometric understanding that allows us to learn a simplicial structure efficiently. Our method requires that the data are unit normalized (and thus lie on the unit sphere). We show that under this restriction, building a model with simplices amounts to constructing a convex hull inside the sphere whose boundary facets is close to the data. We call the boundary facets of the convex hull that are close to the data *Activated Simplices*. While the total number of bases used to build the simplices is a parameter of the model, the dimensions of the individual activated simplices are learned from the data. Simplices can have different dimensions, which facilitates modeling of inhomogeneous data sources. The simplicial structure is bounded — this is appropriate for modeling data with constraints, such as human elbows can not bend more than 180 degrees. The simplices are easy to interpret and extremes within the data can be discovered among the vertices. The method provides good reconstruction and regularization. It supports good nearest neighbor classification and it allows realistic generative models to be constructed. It achieves state-of-the-art results on benchmark datasets, including 3D poses and digits.

1. Introduction

The curse of dimensionality motivates machine learning researchers to search for low-dimensional structures in high-dimensional data. These low-dimensional structures help us understand the data, and they enable us to build data models that avoid over-fitting.

There are several strategies for finding low-dimensional structures. One approach is to project the data into a single low-dimensional linear or non-linear manifold [30, 29, 19, 3, 13] while preserving properties of a local neighborhood graph. These methods are global, in that a single low dimensional structure is used in the representation of the data. One drawback of those methods is that it is difficult to represent some closed manifolds in familiar spaces (such as




Figure 1: (a) A 3-simplex is a solid tetrahedron. Points within the tetrahedron are convex combinations of the 4 vertices $\alpha_1 x_1 + \dots + \alpha_4 x_4$, where $\alpha_i \geq 0$ and $\sum \alpha_i = 1$. (b) A complex consisting of five 2-simplices and some 1-simplices.

Euclidean spaces) of a useful low dimension. These closed manifolds are prevalent in computer vision, for example, in data from repeated human walking motion.

A second approach is to represent data using a combination of simple local structures in the original data space. k -means and its variants represent data by assigning data points to one of k clusters. These methods provide coarse piecewise constant approximations to high-dimensional structures. More recently, k -flats [7, 5], Atlas learning [27], and Sparse Subspace Clustering [18] represent data using a vocabulary of hyperplanes. These methods work very well when the data come from a union of hyperplanes, but are limited when data are from a curved manifold. Since the hyperplanes are unbounded, they don't provide for strict constraints, such as might be appropriate for range of motion restrictions in human pose data.

In this paper we represent data using an arrangement of simplices. Simplices are points, line segments, triangles, tetrahedra, and other higher dimensional analogs of the triangle; see Figure 1. The simplices act as local models of the data manifold. In general, learning a model which represents data using a mixture of simplices is quite difficult. However, we show in Section 3 that, under some reasonable assumptions, the optimal simplices have a convenient geometry, which makes learning dramatically easier.

We assume that the data are normalized to have norm one and hence lie on a unit sphere, and we build the simplices using bases that have norm at most one. With these restrictions, learning the best simplices amounts to finding bound-




Figure 2: Poses as the elbow is articulated. (a) The hand moves along the black curve as the elbow is articulated. The shoulder (blue dot) and elbow (red dot) are fixed in the poses. Our method learns two extreme bases (red and green) and an intermediate basis (cyan). (b) shows two activated simplices (blue line segments). The data on the black curve will be represented by their projections on the simplices. (c) shows a pose sampled from an activated simplex. All poses on the activated simplices satisfy the appropriate range of motion constraints at the elbow.

any facets on the convex hull of the bases that are close to the data. We call the boundary facets activated in this way activated simplices. Once a collection of activated simplices is learned they are pruned to prevent over-fitting.

The Activated Simplices representation has several advantages: (i) Each simplex provides a tight approximation of a local region of the data manifold. (ii) Curved and closed manifolds can be represented using an arrangement of simplices. (iii) The simplices can have different dimensions allowing the representation of an inhomogeneous data source; see Figure 3. While the total number of bases is a parameter of the model, the number and dimensions of the individual simplices are learned. (iv) The representation is bounded. This reduces the risk of generating data which violate constraints in certain applications. See Figure 2. (v) Generative models can be constructed on the simplices using Dirichlet distributions which can be used to synthesize realistic data. (vi) The model is interpretable: points on the simplices are convex combinations of the vertices, and extremes among the data can be found among the vertices. See Figure 7 and 8.

The requirement that data lie on the unit sphere may seem restrictive, however many data sources are normalized in that way, and we present some alternatives to the standard normalization in the Appendix.

The paper is organized as follows. Section 2 reviews related work. Sections 3 and 4 describe the activated simplices learning and pruning methods, and Section 5 discusses optimization strategies. The remainder of the paper is devoted to experiments on digits and poses, followed by a brief conclusion. In the Appendix, we discuss some projections onto the sphere which are useful as alternatives to the usual normalization, we present several low dimensional examples, we discuss the geometry of the standard sparse coding model, and we discuss our method in the context of methods for triangulating manifolds.

2. Related Work

Our method represents data as a mixture of activated simplices. So it is natural to make comparisons with methods which represent data using a mixture of simple structures, such as k -means, k -flats [7] and its generalizations and improvements including Sparse Subspace Clustering [18] and Atlas learning [27]. Since our method builds on a convex hull, it is also natural to compare with Archetypal Analysis [10, 8] which models the data globally as a convex hull and learns convex extremes.

k -means is a clustering algorithm where the data in each Voronoi region is represented by its mean. k -flats extends k -means by using a vocabulary of hyperplanes, usually of the same dimension. The number of hyperplanes and their dimensions are parameters of the model. k -flats is limited in its ability to model curved manifolds, as it must use many hyperplanes for accurate reconstruction.

Sparse Subspace Clustering [18] analyses the *self expression* (i.e., sparse linear interrelations) of the data and groups them according to hyperplanes. Atlas [27] learns hyperplane charts by fitting hyperplanes, with the restriction that an entire local neighbourhood of each point can be represented by the same hyperplane. This improves the assignment of points to hyperplanes at points of ambiguity. These methods give good performance when the data source resembles a small collection of hyperplanes. But they are limited when modelling data from a curved manifold, and the structures learned are not naturally bounded.

Our simplicial model is more flexible in its vocabulary, since it assembles an arrangement of simplices to fit the data. It can manage curved manifolds and inhomogeneous sources (see Figure 3 and the Appendix for some pictures). It naturally accommodates constraints such as anatomical constraints, and allows for generative models by fitting a Dirichlet density to each simplex. It allows classification by nearest neighbour classification on the simplices. It will be preferable to hyperplane methods when the data is not close to a union of hyperplanes.

Archetypal Analysis [10, 8] learns a global convex model for the data. It naturally discovers convex extremes. However when the data manifold is not a convex set, it ignores the details in the interior. The Activated Simplices method constructs local convex structures to describe the data. While it does construct a convex hull to approximate the data on the sphere, it is the boundary simplices of this hull, not the convex hull itself that are important.

3. A Mixture of Simplices Model

In our method we represent data using a mixture of simplices. Simplices are points, line segments, triangles, tetrahedra, and other higher dimensional analogs of the triangle; see Figure 1. A point on a k -dimensional simplex with ver-




Figure 3: The left panel shows the data points from the manifold of a ribbon passing through a circle. The right panel shows the detected basis vectors (red dots) and the activated simplices (line segments in the circle and the small triangles in the ribbon). The simplex dimensions can be different.

tices x_1, \dots, x_{k+1} is a convex combination $\alpha_1 x_1 + \alpha_2 x_2 + \dots + \alpha_{k+1} x_{k+1}$, where $\alpha_i \geq 0$ and $\alpha_1 + \alpha_2 + \dots + \alpha_{k+1} = 1$. A simplicial model built using p bases $\{x_1, \dots, x_p\}$ is a family of simplices $\mathcal{F} = \{\Delta_1, \dots, \Delta_{|\mathcal{F}|}\}$ each of whose vertices is a subset of the bases $X = \{x_1, \dots, x_p\}$. Learning a simplicial model from data needs to solve the complex problem of learning the bases $\{x_1, \dots, x_p\}$, learning the simplices $\mathcal{F} = \{\Delta_1, \dots, \Delta_{|\mathcal{F}|}\}$, learning the assignment of each data point to a simplex, that is, picking a simplex for each data point, and learning the projection of each data point on its chosen simplex. For training data $Y = \{y^{(1)}, \dots, y^{(N)}\}$ this is the following minimization:

$$\min_{X, \mathcal{F}, \alpha, V} \frac{1}{N} \sum_j \sum_{t=1}^N V_t^{(j)} \|y^{(j)} - \Delta_t \alpha_t^{(j)}\|^2 \quad (1)$$

where V is a binary indicator variable assigning each point to exactly one simplex in the family, each $\alpha_t^{(j)}$ is a convex coefficient vector that attempts to express data point $y^{(j)}$ in terms of the t^{th} simplex Δ_t .

We now assume that the data are normalized to have ℓ_2 norm 1 and hence lie on a unit sphere, and that the bases have ℓ_2 norm at most 1, so that they are comparable to the data. This simplifies the assignment problem dramatically, since the data now lie on the surface of a sphere, and the simplices Δ_i are convex combinations of bases within the sphere. It simplifies the problem of identifying optimal simplices, since the convex combinations of the bases that are closest to the data are on boundary facets of the convex hull of the bases. The optimal construction of \mathcal{F} from the bases X is the set of boundary facets activated in this way by the data. Since facets of a convex hull in general positions are simplices, we call these facets activated simplices. The optimal coefficients $\alpha_t^{(j)}$ are those from the projection of the data point $y^{(j)}$ onto the boundary of the convex hull.

We’ve observed that the best projection of $y^{(j)}$ onto the optimal simplices is the projection of $y^{(j)}$ onto the convex




Figure 4: Low dimensional facets are closer to the sphere than high dimensional ones. (a) The interior of the line segment (1-simplex) is further from the circle than its endpoints (0-simplices). (b) The edges of the triangle are closer than its interior.

hull of the bases. This means that in the optimization

$$\min_{\beta^{(j)}} \|y^{(j)} - X\beta^{(j)}\|^2 \quad (2)$$

subject to the constraints $\beta^{(j)} \geq 0$, $\|\beta^{(j)}\|_1 = 1$, which expresses $y^{(j)}$ as a convex combination of **all** the bases, the non-zero entries in the coefficient vector $\beta^{(j)}$ identify the boundary simplex on which $y^{(j)}$ is projected. Thus we can learn a mixture of simplices model by the following minimization

$$\min_{X, \beta} \frac{1}{N} \sum_{j=1}^N \|y^{(j)} - X\beta^{(j)}\|^2 \quad (3)$$

subject to $\beta^{(j)} \geq 0$, $\|\beta^{(j)}\|_1 = 1$, for $j = 1, \dots, N$, and $\|x_i\|_2 \leq 1$ for $i = 1, \dots, p$.

We then derive the activated simplices by studying the activations $\beta^{(j)}$.

3.1. The Activated Simplices Model

To summarize the observations above, we develop a mixture of simplices model for data $Y = \{y^{(1)}, \dots, y^{(N)}\}$ which is assumed to be unit normalized. The simplices are learned from co-activated bases in the optimization (9). The optimization (9) can be understood as learning a convex hull C_X of the bases X , such that its boundary facets are close to the training data. We call the boundary facets that are closest to the data **activated simplices** — these can be determined from activations of the coefficients $\beta^{(j)}$.

The collection of simplices that are learned from (9) can be pruned to obtain a more efficient representation. This is discussed in Section 4.

3.2. Activated Simplices are Low Dimensional

We find in experiments that the activated simplices have low dimension. However the boundary facets for a generic convex body in d -dimensional space are d dimensional simplices. Our method learns facets (that is, sub-simplices of boundary simplices) that have dimension much less than d . We now give an explanation for this phenomenon.




Figure 5: Learning bases using standard sparse coding. (a) Sparse coding represents data using combinations of bases and negative bases. (b) In the usual prior learned from sparse coding, any combination of bases with small ℓ_1 norm is allowed. Here we show combinations with ℓ_1 norm at most one, which corresponds to the convex hull of the bases and negative bases. (c) A sample from this convex hull clearly violates the range or motion constraints at the elbow. The activated simplices in Figure 2 are a better prior.

Suppose that the data Y are from a low dimensional manifold embedded in R^d . Recall that in the minimization (9), a basis X is learned so that boundary facets of the convex hull C_X are close to the data. The learning process constructs X so that this representation is as efficient as possible.

The interaction of the curved geometry of the sphere with the linear geometry of the boundary facets of C_X , means that it is most efficient to position low dimensional facets close to the data. Figure 4 shows a segment (a 1-simplex) with its vertices in the circle, and it shows a triangle (a 2-simplex) with its vertices in the 2-sphere. Notice that the interior of the segment is further from the sphere than its endpoints, similarly the interior of the triangle is further from the sphere than its boundary segments. In general, it is more efficient to approximate points on the sphere locally with low dimensional facets, than with high dimensional facets. Hence the optimization (9) learns a structure that approximates the data with low dimensional facets of C_X .

Of course if the data is distributed uniformly over the sphere, it will be impossible to position low dimensional facets close to the data points, without a vast number of bases. But we see in our experiments, with what are understood to be low dimensional data sources, that the activated simplices are low dimensional. This makes sense as it provides a most efficient representation of the data.

3.3. Relation to Sparse Coding, and a Penalized Version of Activated Simplices

The coefficients vectors $\beta^{(j)}$ learned in (9) are sparse, that is, most entries are zero, and the non-zero entries identify a low dimensional boundary facet of the convex hull of the bases X . Some readers will be reminded of sparse coding. Indeed it is possible to understand sparse coding using similar ideas — we discuss this in the Appendix.

In the usual formulation of sparse coding [31, 26, 4, 20]

data Y are represented by bases X through the following ℓ_1 penalized optimization:

$$\min_{X, \beta} \frac{1}{N} \sum_{j=1}^N \|y^{(j)} - X\beta^{(j)}\|^2 + \lambda \|\beta^{(j)}\|_1 \quad (4)$$

The bases x_i are constrained to have ℓ_2 norm at most 1. Notice there is no constraint here that the coefficients $\beta^{(j)}$ are positive. The method learns sparse representations of the data as combinations of the bases and negative bases.

The bases learned in sparse coding can be used as a regularizer for pose data, in a similar spirit to our experiments. See [32]. However the usual prior associated with a sparse coding representation, is that any combination $X\beta$ can occur, provided that $\|\beta\|_1$ is small. This allows co-activations of the bases that don't occur in the training data. Figure 5 shows a regularizing structure that might be learned in this way. Compare with Figure 2, to see that the Activated Simplices method learns a more realistic representation.

It is interesting to experiment with a penalized version of the Activated Simplex method. A formulation with positive penalty parameter $r \leq 1$ is

$$\min_{X, \beta} \frac{1}{N} \sum_{j=1}^N \|y^{(j)} - X\beta^{(j)}\|^2 \quad (5)$$

subject to $\beta^{(j)} \geq 0$, $\|\beta^{(j)}\|_1 = r$, for $j = 1, \dots, N$, and $\|x_i\|_2 \leq 1$ for $i = 1, \dots, p$.

This will force representations on lower dimensional facets when $r < 1$. We've experimented with this penalized version and found that the unpenalized version works well with the data we considered. But we can imagine that the additional sparsity provided in (13) might be useful in managing some high dimensional data.

4. Pruning The Activated Simplices

Each training point $y^{(j)}$ proposes (activates) a simplex Δ . Gathering the activated simplices, we obtain a set of simplex proposals $\mathcal{F} = \{\Delta_1, \dots, \Delta_T\}$. The number of proposed simplices will typically be much less than the number of training data, since several data points will activate the same simplex.

To prevent over-fitting, we propose a principled way to prune the simplices \mathcal{F} to a proper subset \mathcal{F}^* without degrading the reconstruction results much. For example, \mathcal{F} might contain a high dimensional simplex Δ that received few activations, but its low dimensional sub-simplices might have received many activations. We may obtain a better model if we remove this large simplex from the model.

We prune so as to obtain a small number of simplices and so that the remaining simplices have low dimensions. We use the following to balance the goals of few and low




Figure 6: (a) The effect of simplices pruning. The x-axis is the number of the selected simplices. The y-axis is the corresponding reconstruction error. The number of pre-pruned simplices is about 3,000. The method obtains comparable reconstruction errors by selecting about 30 simplices from the 3,000. (b,c,d) show the influence of the two parameters λ_1 and λ_2 on the reconstruction error, the number of simplices and the average dimensions of the simplices.




Figure 7: The top row shows sample faces of different lighting directions in terms of Azimuth and Elevation in the Yale Face B dataset. The mid row shows the learned bases. We can see that some of them approximately correspond to the faces of extreme lighting conditions. The bottom row shows the reconstructed faces and residuals.

dimensional with the goal of good reconstruction:

$$\mathcal{F}^* = \underset{\hat{\mathcal{F}} \subset \mathcal{F}}{\operatorname{argmin}} L(Y, \hat{\mathcal{F}}) + \lambda_1 \#(\hat{\mathcal{F}}) + \lambda_2 \sum_{i=1}^{|\hat{\mathcal{F}}|} \dim(\hat{\mathcal{F}}_i), \quad (6)$$

where $L(Y, \hat{\mathcal{F}})$ represents the reconstruction error using the set $\hat{\mathcal{F}}$ of simplices to approximate the data, and $\#(\hat{\mathcal{F}})$ is the number of simplices in the set $\hat{\mathcal{F}}$, and $\dim(\hat{\mathcal{F}}_i)$ is the dimension of the i^{th} simplex in the set $\hat{\mathcal{F}}$. The two penalty terms $\lambda_1 \#(\hat{\mathcal{F}})$ and $\lambda_2 \sum_{i=1}^{|\hat{\mathcal{F}}|} \dim(\hat{\mathcal{F}}_i)$ encourage fewer simplices and encourage lower dimensions. We set λ_1 and λ_2 to be $0.001/T$ and $0.01/T$, respectively, for our experiments. In general, these can be set by cross-validation. Figure 6 shows the influence of pruning.

5. Optimization

The formulation (9) is non-convex, but it is convex with respect to each of the variables, X and β , when the other is




Figure 8: The figure shows 30 human poses from a walking sequence (in green) and the learned three bases (in red) overlaid together, viewed from two different directions. We can see that our method actually learns three extreme bases.

fixed. Optimizing the two convex sub-problems alternately leads to a local minimum, but using classic convex optimization methods may be slow when the number of training data is large. Inspired by [23], we solve the problem in an online way, based on stochastic gradient descent. This scales up gracefully to large dataset. When we fix β and optimize the bases X , we obtain the same dictionary update problem as in [23]. So we refer the readers to [23] for more information. When we fix X and update β , we obtain a least-squares problem with a simplex constraint. We use an active-set algorithm [35] to benefit from the sparsity of β .

The problem (6) is difficult to optimize directly because of the large search space. Hence we use a greedy approach which successively adds simplices to $\hat{\mathcal{F}}$ which most reduce the objective (6). We terminate the process when the objective stops decreasing. On a regular Desktop equipped with 3.4GHz CPU, the overall simplices learning time is about 0.19 seconds for a digit dataset containing 800 training data of dimension 256. Projecting data on the simplices is fast.

6. Data Reconstruction

In this section we evaluate the reconstruction errors by computing the Euclidean distance between data and their projections on the nearest simplices. We evaluate the influence of the model parameters p , i.e., the number of bases. We also show some results when the convex hull radius r is




Figure 9: Reconstruction errors in terms of the number of bases and the scale r of the convex hull. The left figure shows the reconstruction error when training and testing on the same actions in the dataset. We can see that the least reconstruction error (achieved when $r = 1$ and $p = 500$) is almost zero which justifies the activated simplices approximate the data well. The right figure shows the reconstruction errors when training and testing on different actions.

smaller than one.

We conduct experiments on a large human pose dataset H3.6M [21] which contains 3.6 million poses. We select 11,000 poses of 11 actions including “walking”, “taking photo”, “smoking”, “sitting down”, “sitting”, “purchases”, “posing”, “phoning”, “greeting”, “eating” and “discussion”.

We first evaluate on the training data—we train and test on the same 11,000 poses. We learn a single set of simplices for the eleven actions together. Figure 9 (left) shows the average reconstruction errors for each choice of p and r . The smallest average reconstruction error is about 0.0423 (achieved when $r = 1$ and $p = 500$) which indicates that the activated simplices can represent the training data well. The reconstruction errors decrease fast as radius r increases. The reconstruction errors also decrease as the number of bases p increases but in a more stable manner.

In the second scenario, we split the 11,000 poses into training and testing subsets each containing 5,500 poses of the 11 actions. We learn the activated simplices from the training data and test on the testing data. Figure 9 (right) shows the results. We can see that the reconstruction error is similar to that of the first experiment. The smallest reconstruction error is about 0.05 (achieved when $r = 1$ and $p = 500$). The results justify that the activated simplices model can reconstruct the data well. Figure 10 shows three sample reconstruction results whose reconstruction errors are 0.06, 0.1 and 0.2, respectively.

7. Nearest Neighbour Classification

We conduct digit classification on the Semeion handwritten digit dataset [1] and action classification on a more challenging MSR-Action3D dataset [22]. We use simple nearest neighbour classifiers on the simplicial model and achieve the state-of-the-art performance on both tasks.




Figure 10: Sample reconstruction results. The poses in black are the groundtruth poses and the red ones are their projections on the activated simplices. The reconstruction errors are 0.06, 0.1 and 0.2, respectively, from left to right.

Table 1: Digit classification accuracy (%) of sparse coding, LTSA [36], Atlas [27], Archetypal Analysis [8] and our approach. The best performance of different dimension choices is reported for LTSA and Atlas. The dimension of our method is automatically determined by the algorithm.

Methods	Best Dimensionality	Accuracy(%)
Sparse Coding	10	89.40
LTSA [36]	34	90.44
Atlas [27]	31	91.73
Archetypes [8]	10	91.17
Our	4-8	93.00

7.1. Handwritten Digit Classification

The Semeion dataset [1] contains 1593 grey scale images of handwritten digits (0, . . . , 9) generated from about 80 people who each wrote all ten digits twice. The images are of size 16×16 and each pixel in the image is converted into a boolean value using a fixed threshold. We randomly select 796 images for training and 797 images for testing following [28]. The process is repeated for 100 times and the average result is reported. We compare with the standard manifold learning methods such as LTSA [36], Atlas [27], Archetypes [8] and the usual sparse coding.

In the training stage, for our method, we learn a set of activated simplices for each of the ten classes using the same parameters. In the testing stage, for each data, we project it to the ten classes of simplices independently and obtain the class having the smallest reconstruction error. We set the number of bases p to be ten for each class (100 in total). We finally obtain 100 simplices whose dimensions are from four to eight. For sparse coding and Archetypes methods, we learn a set of bases for each of the ten classes and project test data to the ten sets of bases independently and obtain the class which has the smallest error.

Table 1 shows the recognition results. We outperform all of the methods. Our approach automatically determines the simplex dimensions which vary from simplex to simplex. For other methods, different choices of dimensions have been tried and the one that achieves the best performance is reported. The number of bases is 100 for sparse




Figure 11: Classification accuracy on the Semeion dataset using different numbers of bases and the convex hull radius. In the left figure, the classification accuracy begins to decrease when the number of bases is larger than 400. This is because too large a number of base causes over-fitting.

Table 2: Action recognition accuracy on MSR-Action3D.

Methods	Accuracy (%)
Sparse Coding	80.22
Action Graph on Bag-of-3D joints [22]	74.70
Actionlet Ensemble [34]	88.20
Spatial-Temporal-Part model [33]	90.22
Our Approach	91.30

coding and archetype for fair comparison. Figure 11 shows the influence of the number of bases and the radius of the convex hull on the classification results.

7.2. Action Recognition

MSR-Action3D dataset [22] provides 557 human pose sequences of ten subjects performing 20 actions which are recorded with a depth sensor. There are about 50 frames in each sequence. This is a challenging dataset because many of the actions are highly similar to each other. We use the cross-subject evaluation scheme as in [22].

To make the poses more discriminative between similar classes, for each pose in the sequence, we stack it with its consecutive ten poses as a high dimensional pose snippet. Weak temporal coherence can be achieved by this simple stacking operation. The activated simplices are learned based on snippets rather than a single pose.

In the training stage, for each of the 20 classes, we learn an independent set of activated simplices using the same parameters. In the testing stage, for each sequence, we project the 3D snippets onto the 20 sets of activated simplices and select the class which has the smallest reconstruction error. We set the number of bases for each class to be 25 (500 in total) by cross-validation.

Table 2 shows the results. Our method outperforms [22] and sparse coding, and achieves comparable performance as [34] and [33]. However, actionlet ensemble uses sophisticated feature learning methods and multiple kernel learning classifiers. Spatial-temporal-part model uses data mining techniques to remove the noisy joints and then use support vector machine for classification. In contrast, our method is




Figure 12: Left figure compares the estimation results of sparse coding, archetypal analysis and our method. The x-axis is the reconstruction error. The y-axis is the percentage of data whose estimation errors are less equal than the error specified by the x-axis. The right figure shows the average limb length of the estimated 3D poses of the three methods and ground-truth. The limb lengths of our method are much closer to the ground-truth which shows that the model better respects the anthropomorphic constraints.

the simplest in terms of both features and classifiers.

We have an interesting observation from the experiment. For each test sequence, if we select the top K classes which have the smallest reconstruction errors and regard the classification is correct when any of the K predictions is correct, then we can considerably improve the performance. For example, when $K = 2$, the accuracy increases to 95.54%; when $K = 5$, the accuracy increases to 97.03%; when $K = 8$, the accuracy is 100%.

8. 3D Human Pose Estimation

In this section, we apply the simplicial model to the task of 3D human pose estimation from 2D poses. This allows us to compare our method with the more standard application of sparse coding [32].

The 3D and 2D poses are represented by n joint locations $P \in \mathbb{R}^{3 \times n}$ and $O \in \mathbb{R}^{2 \times n}$, respectively. We conduct experiments on the H3.6M Human motion capture dataset [21] which provides ground truth 3D human joint locations of 11 actions. To evaluate the generalization ability of the method, we select the poses from the first five actions for training and the remaining six actions for testing.

We obtain the 2D joint locations by projecting the 3D pose into 2D using synthetic weak perspective camera parameters M , i.e., $O = M \cdot P$, where $M = \begin{pmatrix} m_1^T \\ m_2^T \end{pmatrix} \in \mathbb{R}^{2 \times 3}$. We assume the joint locations are mean-centered hence eliminate the translation component for simplicity. We now have 3D joint locations P which we use as ground truth, and we have the corresponding 2D joint locations O — our goal is to reconstruct P from O with the help of the activated simplices.

We represent the 3D poses by activated simplices. We first learn a set of activated simplices $\mathcal{S} = \{\mathcal{S}_1, \dots, \mathcal{S}_K\}$ from the training 3D poses. Then for a 2D pose O , we op-

timize for the simplex $\mathcal{S}_* \in \mathbb{S}$ and the simplex coefficients β_* by minimizing the following:

$$\begin{aligned} \mathcal{S}_*, \beta_* = \operatorname{argmin}_{\mathcal{S} \in \mathbb{S}, \beta} \quad & \|M\mathcal{S}\beta - O\|_2^2 \\ \text{s.t.} \quad & \beta \geq 0, \quad \|\beta\|_1 = 1 \end{aligned} \quad (7)$$

This minimization can be done naively since the number of simplices is small. Then we reconstruct its 3D pose P by $\hat{P} = \mathcal{S}_*\beta_*$. The reconstruction error is the average Euclidean distance between each joint of P and \hat{P} .

We compare our method with the classical sparse coding [23] and archetypal analysis [10] methods. For these two methods, we first learn a set of bases or archetypes X on the training data and then for each test data, we optimize the basis coefficients β_* by minimizing the error between the 3D pose projection and the 2D pose:

$$\beta_* = \operatorname{argmin}_{\beta} \|MX\beta - O\|_2^2. \quad (8)$$

For Archetypal method we constrain β is non-negative and sum to one. The estimated 3D pose is $\hat{P} = X\beta_*$.

Figure 12 (left) shows the human pose estimation results. Our method achieves better performance than the other two alternatives. There could be two reasons accounting for this: (1) compared with sparse coding, the activated simplices constrain the combinations of bases to combinations that occur in reality, whereas the restrictions in sparse coding methods are looser. Figure 12 (right) shows that the estimated limb lengths of our method are close to the ground-truth while those of the sparse coding and archetypes are considerably different from the ground-truth. The experimental results show that the sparse representation cannot implicitly enforce the limb length constraint; (2) compared with archetypal analysis method, the activated simplices explore the local structures including the interior of the manifold which makes the method more accurate. In contrast, archetype can only learn the convex hull of the manifold and fail to explore the interior.

9. 3D Human Pose Synthesis

Our method also allows to synthesize realistic data. For each simplex \mathcal{S}_i , we learn a Dirichlet distribution [24] from the coefficients $\beta_i^j, j = 1, \dots, m$ of the m data that are projected to the simplex. To generate a data, we first sample a simplex and then sample the coefficients β from the Dirichlet distribution of that simplex. We output the corresponding combination of bases as the synthesized data.

Dirichlet distribution is often used as prior distributions in Bayesian statistics. The probability density function is defined as: $P(\beta) = \frac{1}{B(\alpha)} \prod_{i=1}^m \beta_i^{\alpha_i-1}$, where $\beta > 0$ and $\|\beta\|_1 = 1$. $B(\alpha)$ is a normalization factor. We learn the parameters α by Maximum Likelihood Estimation.




Figure 13: Synthesized human poses (in green) are overlaid with their nearest neighbours (in black) in the training dataset. The actions of the nearest neighbours are shown in the texts below the poses. We can see that the synthesized poses differ from the training data but they are still realistic which shows the simplices' generalization properties.




Figure 14: Left figure shows the limb lengths of the synthesized poses and of the training data (groundtruth). The right figure shows the range of bending angles of four body joints in the training data (green bars) and the bending angles of the synthesized data (blue dots). We can see that the bending angles of the synthesized data are within the limit of those of the training data.

We conduct experiments on the H3.6M human pose dataset. We select 55,000 poses from 11 actions and learn a single set of simplices. We set the number of bases to be 100. The algorithm ends up with 83 simplices. We learn Dirichlet distributions on each simplex. We display some typical synthesized poses in Figure 13. We can see that they are divergent and realistic. In addition, Figure 14 shows that the synthesized poses also respect well the limb length and bending angle constraints.

10. Conclusion

We propose a method for representing data using a mixture of activated simplices. The Activated Simplices representation allows accurate reconstruction by preventing combinations of bases that do not occur in the data.

Acknowledgement: The project is supported by US Army Research Office ARO Proposal 62250-CS and NSF STC award CCF-1231216 and Center for Minds, Brains and Machines (CBMM)

Appendix starts here

11. Introduction

In the appendix we provide more details on the following:

- (i) The relationship between Activated Simplices and traditional sparse coding.
- (ii) Alternatives to standard normalization for projection onto the sphere.
- (iii) Triangulating Manifolds.
- (iv) Extremes in data and a comparison with Archetypal Analysis.

12. Activated Simplices and Sparse Coding

In Section 3 of the Main Paper we discussed how a mixture of simplices model can be learned by constructing a convex hull inside the sphere whose boundary facets are close to the data. This will remind readers of traditional sparse coding [25], since the projections of the training data on this hull are sparse, i.e., they involve relatively few of the bases. We now explore this connection in more detail.

Our model requires normalized data $Y = \{y^{(1)}, \dots, y^{(N)}\}$ and it constructs a convex hull of p bases $X = \{x_1, \dots, x_p\}$ such that boundary facets are close to the data. It constructs this hull by the following minimization

$$\min_{X, \beta} \frac{1}{N} \sum_{j=1}^N \|y^{(j)} - X\beta^{(j)}\|^2 \quad (9)$$

subject to $\beta^{(j)} \geq 0$, $\|\beta^{(j)}\|_1 = 1$, for $j = 1, \dots, N$, and $\|x_i\|_2 \leq 1$ for $i = 1, \dots, p$.

The activated simplices correspond to co-activated bases.

Figure 15 shows a very simple situation where a convex hull is constructed from data on the circle using 6 bases. The activated simplices are the three boundary segments closest to the data.

In the usual formulation of sparse coding, a dictionary of p bases $X = \{x_1, \dots, x_p\}$ is constructed to represent data $Y = \{y^{(1)}, \dots, y^{(N)}\}$. The dictionary is constructed by the following minimization

$$\min_{X, \beta} \frac{1}{N} \sum_{i=1}^N \left(\frac{1}{2} \|y^{(i)} - X\beta^{(i)}\|^2 + \lambda \|\beta^{(i)}\|_1 \right) \quad (10)$$

subject to $\|x_i\|_2 \leq 1$ for $i = 1, \dots, p$. Here, λ is a non-negative penalty parameter.

Notice some differences between (9) and (10). In (9) the coefficients $\beta^{(j)}$ are required to be non-negative, but there




Figure 15: A simple picture showing data on the circle (red dots) and a convex hull fitted to the data using 6 bases (blue dots).

is no such constraint in (10). In (9) the coefficients $\beta^{(j)}$ are constrained to have ℓ_1 norm 1, whereas in (10) the ℓ_1 norms of the coefficients are penalized through the $\lambda \|\beta^{(j)}\|_1$ term. But both formulations have a similar quadratic data reconstruction term $\|y^{(j)} - X\beta^{(j)}\|_2^2$.

It is possible to interpret this sparse coding in geometrical terms [12, 20], and it is likely that this geometrical understanding was behind Tibshirani's original formulation of the lasso [31].

To begin, we describe the geometry of a single penalized regression. In this situation the bases X are known, and we are interested in the coefficients β for a single data point y . This corresponds to the minimization [31, 26]

$$\min_{\beta} (\|y - X\beta\|^2 + 2\lambda \|\beta\|_1) \quad (11)$$

In this regression the penalty parameter λ and the data point y control the ℓ_1 norm of the coefficients, that is, they determine a *radius* $r = r(y, \lambda)$ such that the coefficients β solve

$$\min_{\beta} \|y - X\beta\|_2 \quad (12)$$

subject to $\|\beta\|_1 \leq r$. As the penalty parameter λ increases, the radius r decreases.

We can interpret this minimization in terms of convex geometry. The set of $X\beta$ such that $\|\beta\|_1 \leq r$ is the convex hull of the bases $X = \{x_1, \dots, x_p\}$ and the negative bases $-X = \{-x_1, \dots, -x_p\}$, scaled by a factor of r . Figure 16 shows bases, and negative bases, and their convex hull, scaled by r .

For the regression of a single point, the penalty λ determines a convex hull of radius r , and then the regression coefficients β correspond to the closest point on the boundary of the scaled convex hull. See Equation 12.

The geometry is more complicated when more than one point is regressed on X , as the radius r depends on the penalty λ and on the data point y . Figure 17 shows that r varies with y . In this figure points on the circle are regressed on 3 bases with $\lambda = 0.15$ and $\lambda = 0.6$. The blue curve shows how $r(y)$ varies as y moves along the circle.




Figure 16: The convex hull of bases and negative bases scaled by r .




Figure 17: The radius r (blue curve) in penalized regression depends on λ and y . On left $\lambda = 0.15$, on right $\lambda = 0.6$.

This gives us a geometrical understanding (albeit a somewhat complicated one) of the learning process in the usual formulation of sparse coding. The method positions the convex hull of the bases and negative bases, so that boundary facets of some scalings of this hull are close to the data. The complication here is that the scalings are not uniform, they depend on the $y^{(j)}$.

Our method (9) makes two major changes to the sparse coding formulation. First, the radius r is set uniformly to 1 for all data points $y^{(j)}$, thus we can think of the data as projecting on the same convex hull. Second, the coefficients $\beta^{(j)}$ are restricted to be non-negative, thus we consider the convex hull of the bases, rather than the hull of the bases and negative bases.

We might have learned co-activations from the original sparse coding formulation, but the modifications we've made improve performance as well as interpretability. The uniform radius requirement $r = 1$ encourages the fitting of low-dimensional facets to the data. This encouragement is weaker in the usual sparse coding. Figure 18 shows the projection (blue curves) of the circle learned using the usual sparse coding with three bases and their negatives. Notice that the projections are no longer polygonal. This fitting of a bulging polytope lessens the incentive to fit low dimensional facets close to the data. The left panel in Figure 18 uses $\lambda = 0.25$ and the right panel used $\lambda = 0.6$.

The second modification restricts the coefficients $\beta^{(j)}$ to be non-negative. Without this restriction, a basis might be used with a positive coefficient in some combination to approximate data in some region of the data manifold and it




Figure 18: The projection (blue curve) of the circle learned using 3 bases and their negatives. Notice the bulging. The left panel uses $\lambda = 0.25$ and the right uses $\lambda = 0.6$.




Figure 19: Decreasing r encourages sparsity. Points on the orange cones project to 0-simplices.

might be used with a negative coefficient to represent data somewhere else, far off on the data manifold. The positive restriction means that bases are used in a more focused way in reconstruction and learning.

12.1. Penalized Activated Simplices

We presented a penalized version of our activated simplices formulation in Section 3.3 of the Main Paper. We might have applied an ℓ_1 penalty as in the usual sparse coding (10), but we prefer to control the geometry uniformly and explicitly through r . The penalized formulation with penalty parameter $r \leq 1$ is

$$\min_{X, \beta} \frac{1}{N} \sum_{j=1}^N \|y^{(j)} - X\beta^{(j)}\|^2 \quad (13)$$

subject to $\beta^{(j)} \geq 0$, $\|\beta^{(j)}\|_1 = r$, for $j = 1, \dots, N$, and $\|x_i\|_2 \leq 1$ for $i = 1, \dots, p$.

This can be understood geometrically as learning a convex hull of radius r , such that boundary facets are close to the data. Figure 19 shows how decreasing r encourages sparsity. In the figure, points in the orange regions project onto the vertices, and as r decreases more of space is covered by these orange regions.

This penalization will necessarily worsen reconstruction, but it is useful if one is interested in dimension reduction. In our experiments, which focused on reconstruction from real low dimensional data sources, we found that $r = 1$ was best. See Figure 9 and 11 of the Main Paper where we




Figure 20: The torus is flattened into a band when normalized.

experiment with the number of bases p and the radius r for reconstruction of poses and digits. We experiment with r in reconstructions of the torus from synthetic data in Section 13.

13. Alternative Projections

The Activated Simplices method requires that training data lie on the sphere. This is usually accomplished by directly normalizing the data, or by centering and then normalizing.

The standard normalization is the transformation

$$y^{(j)} \mapsto \frac{y^{(j)}}{\|y^{(j)}\|_2} \quad (14)$$

and centering and normalizing is the transformation

$$y^{(j)} \mapsto \frac{y^{(j)} - \bar{y}}{\|y^{(j)} - \bar{y}\|_2} \quad (15)$$

where \bar{y} is the data mean $\frac{1}{N} \sum_{j=1}^N y^{(j)}$.

For many signals little information is lost in replacing data by its normalization. Indeed this normalization is a common first step in many analyses, for example, images are often contrast normalized. In the experiments in this paper the data was centered and normalized.

However some structure may collapse under direct normalization. Consider, for example, the usual 2 dimensional torus. Figure 20 shows the torus and its normalized image in the sphere — the normalization flattens the torus into a band around the equator and the inner cavity is lost.

To avoid flattenings such as this we use *stereographic projection* to map data in \mathbb{R}^d into the d dimensional sphere \mathbb{S}^d in \mathbb{R}^{d+1} . This is a standard technique in map making where our world (the 2-dimensional sphere) is mapped onto a flat sheet (\mathbb{R}^2). Figure 21 shows stereographic projection from the equatorial plane onto the sphere, through the north pole; the point P maps to the point Q . The point Q is the point on the line through N and P that lies on the sphere.

Stereographic projection maps \mathbb{R}^d onto the d -sphere with the north pole removed (the points at infinity map to the




Figure 21: Stereographic projection from the equatorial plane into the sphere maps P to Q .




Figure 22: A simplex in the plane and the corresponding curved simplex on the sphere.

north pole.) It is an invertible conformal map (it preserves angles between curves). The map is

$$\pi : P \mapsto Q = N + \frac{2}{1 + \|P\|^2} (P - N) \quad (16)$$

and the inverse is

$$\pi^{-1} : Q \mapsto P = N + \frac{1}{1 - Q_{d+1}} (Q - N) \quad (17)$$

where Q_{d+1} is the $d + 1$ st coordinate of Q .

In some situations the entire data analysis can be carried out in \mathbb{S}^d rather than in \mathbb{R}^d . For a classification task, one might map the data to the sphere with stereographic projection, learn a simplicial structure, and carry out classification tasks there. For synthesis one might map the synthesized points back to \mathbb{R}^d with inverse stereographic projection.

But it can also be useful to map the simplicial structure back to \mathbb{R}^d . This is simply a matter of keeping track of activations. The simplex in \mathbb{R}^d interpolating $\pi^{-1}(x_{i_0}), \dots, \pi^{-1}(x_{i_k})$ maps to the curved simplex in \mathbb{S}^d with vertices x_{i_1}, \dots, x_{i_k} , which lies over the flat activated simplex with vertices x_{i_1}, \dots, x_{i_k} on the convex hull! Figure 22 illustrates the relationship between a simplex in the plane and a curved simplex on the sphere.

This relationship between a simplicial structure in Euclidean space and a convex structure in the sphere will remind some readers of [6], where it is revealed that the Delaunay triangulation for some data can be constructed by mapping the data into the sphere with stereographic projection, finding the convex hull of the projected points, and




Figure 23: torus using $N = 1000, p = 60, r = 1$




Figure 24: hemisphere using $N = 1000, p = 50, r = 1$

mapping the faces of this hull back to Euclidean space. Our method approximates the convex hull of the projected points and builds a simplicial approximation to the data in Euclidean space by mapping some facets of this hull back to Euclidean space. It is also interesting to compare with [17] which shows that the Delaunay triangulation can be found from a convex hull after mapping the data into a paraboloid. The common theme is that mapping data into the boundary hyper-surface of a convex region allows useful triangulations to be constructed by convex approximation.

We'll say a little more about Delaunay triangulations in Section 14.

13.1. Low Dimensional Examples.

We now discuss some experiments with some data from some low dimensional manifolds. We sample data from some 2 dimensional manifolds embedded in \mathbb{R}^3 ; we map the data into the 3-dimensional sphere \mathbb{S}^3 using stereographic projection and build a simplicial structure; we map the simplicial structure back into \mathbb{R}^3 and show some plots. No pruning was used here.

Figure 23 shows a structure learned from $N = 1000$ data points sampled from the torus, with $p = 60$ bases using $r = 1$ (unpenalized), and with no pruning. Some 3-dimensional simplices were activated but only their 2-dimensional boundaries are shown here.

Figure 24 shows a simplicial structure learned from 1000





Figure 25: Swissroll and plane using $N = 2000, p = 100, r = 1$ viewed from two directions.

points sampled from the hemisphere, using 50 bases and $r = 1$. No simplices of dimension more than 2 were activated and there are no holes in the structure. This is an excellent reconstruction of the hemisphere.

Figure 25 shows a structure learned from 2000 points sampled on the union of a ribbon and a planar segment, using 100 bases and $r = 1$. A small number of 3-dimensional simplices were learned.

Figure 36 shows simplicial structures learned from the torus, using 1000 samples, and for several values of p and r . Notice that as r gets smaller the activations become more sparse, and for small values of r the structure is a wireframe. A heavily penalized structure constructed using 5 bases and $r = 0.4$ and $r = 0.1$ captures the medial axis of the torus.

14. Triangulating Manifolds

In the main paper we focused on Activated Simplices as a method for data reconstruction, regularization and synthesis. We now discuss Activated Simplices as a tool for manifold reconstruction. This hasn't been a priority in this work, but it is a natural extension. The illustrations in Section 13.1 show the potential for reconstructing low dimensional manifolds. Figure (25) shows a reconstruction of a structure with intersecting components.

It is useful to distinguish the task to building a struc-




Figure 26: A sphere and a wireframe. The wireframe is good enough for data regularization, but it is a poor reconstruction of the manifold.

ture for data reconstruction and regularization from that of manifold reconstruction. Figure 26 shows a sphere and a wireframe. A fine wireframe is a good enough structure to regularize data from the sphere but it is a poor reconstruction of the manifold. The goal in manifold reconstruction is usually to construct a smooth manifold, or a union of a small number of smooth manifolds, that approximate a data source.

The pruning process in Section 4 of the Main Paper is aimed at economical reconstruction, but it may be inappropriate for manifold reconstruction where smoothness is a priority. The plots in Section 13.1 show unpruned structures.

It is interesting to compare Activated Simplices with methods which construct simplicial structures for manifold reconstruction, especially those that use the Delaunay complex [16, 14, 2, 15, 11, 9]. The Delaunay complex for a set of data points is a simplicial structure that is dual to the Voronoi cells for the points. For points in general position in the plane, three points are connected by a Delaunay triangle if their Voronoi regions meet at a point. Figure 27 shows Voronoi boundaries and the corresponding Delaunay triangles for 6 points in the plane. Figure 28 shows the Voronoi cells and Delaunay triangles for 50 points sampled on an ellipse. The Delaunay complex in d dimensions is defined analogously: $d + 1$ points are connected by a d -dimensional simplex if their Voronoi regions intersect.

The Delaunay complex for points in \mathbb{R}^d is a d dimensional structure. It must be pared back to a lower dimensional structure to reconstruct a lower dimensional manifold. The methods that use Delaunay complexes differ in how this paring back is done. For example, α -shapes [16] extracts sub-simplices of bounded diameter, whereas co-cone methods [2] select simplices by their proportions.

Activated Simplices doesn't construct a sub-complex of the Delaunay complex, but it is related to Delaunay complex methods in that it learns from the convex hull of the data in the sphere — [17] shows that the Delaunay complex can be constructed from the convex hull of the stereographic projection of data onto the sphere. The main difference between Activated Simplices and methods which pare




Figure 27: Some points (black dots), the Voronoi boundaries (red) and the Delaunay triangulation (blue)




Figure 28: Top: the Voronoi cells for some points sampled on an ellipse. Bottom: the corresponding Delaunay triangulation.




Figure 29: (a) Boundary points have a different neighborhood structure: q is a boundary point on the closed square, but p is not. (b) The blue circle is the boundary of the bowl.




Figure 30: Left: a car tire (not the inner tube, but the tire). Middle: the convex extremes of the manifold is the set of points where the tire might meet the road. Right: the two blue circles are the boundary points.

back the Delaunay complex is that the simplices are selected by Activated Simplices because they are activated by data, whereas in the Delaunay complex methods, sub-simplices of the complex are selected if they meet some geometric requirements such as size or proportion.

15. Extremes and Archetypes

We now discuss how extremes within data might be identified with Activated Simplices. But first we must define some notions of extreme. A first notion is what might be called a *convex extreme point*. This is a point on the data manifold that is not a convex combination of any other points on the manifold. For example, in a cube, the corner points are convex extremes because they are not convex combinations of other points. On a 2-dimensional sphere, every point is a convex extreme point. This is an extrinsic notion of extreme — a person looking on the manifold from afar can identify convex extremes, but a short sighted ant crawling on the manifold might notice nothing unusual as it crawls past a convex extreme. Figure 30 shows a tire and its convex extremes. The convex extremes of a finite set of points are the vertices on the convex hull of the points.

Archetypal Analysis [10] [8] learns a convex model for the data (it learns bases such that the convex hull encloses the data). The bases that are learned are approximations to extreme points.

A second notion of extreme point comes from the mathematical notion of a manifold with boundary. In this formulation an interior point has neighborhoods homeomorphic to disks, but boundary points have neighborhoods homeomorphic to half disks. Figure 29 (a) shows a neighborhood for an interior point p and a boundary point q on the square. Figure 29 (b) shows the boundary points of the bowl and Figure 30 shows the boundary points of the tire.

Boundary point is an intrinsic notion — a short sighted ant will notice that the neighborhood is different at a boundary point. Boundary points correspond to extreme configurations. For example, a pose, where a limit of the range of motion at some joint is reached, is a boundary point on the manifold of poses.

Figure 30 illustrates that convex extremes and boundary points may not coincide.

15.1. Boundaries and Activated Simplices

We now develop a notion of boundary for a simplicial structure. Figure 31 shows the bowl and part of a triangulation. The blue segments are boundary simplices. Figure 32 shows another simplicial structure; the blue simplices are boundary simplices, the red segment is an internal segment. These pictures suggest a definition of boundary simplex — a simplex is a boundary simplex if it has co-dimension 1 in a maximal simplex and it is not a sub-simplex of two maximal simplices. In Figure 32 the blue segments are sub-simplices of only one maximal simplex, whereas the red segment is a sub-simplex of two maximal simplices.

We can use this notion to detect extremes using Activated Simplices. Figure 33 shows some boundary faces found in the Yale B face data. We used 64 face images of the same person under different lighting conditions. We




Figure 31: The bowl and the top part of a triangulation, with boundary simplices in blue.

learn 10 bases, obtain 27 simplices and identify 14 maximal simplices. Seven of the maximal simplices are in the interior of the manifold. It means all of their co-dimension 1 sub-simplices are sub-simplices of other maximal simplices. Each of the remaining seven maximal simplices has one co-dimension 1 sub-simplex that is a boundary. Figure 33 shows vertices for the seven boundary simplices. We can see that the boundary simplices include the images corresponding to extreme lighting directions.

The utility of the boundaries found from Activated Simplices depends on the data manifold and on the simplicial structure that is constructed. If the simplicial structure is stringy and lumpy there will be many boundary simplices, and it will be hard to make sense of these extremes. This kind of simplicial structure can arise because it reflects the reality of the data manifold, or it can occur when the data manifold is under-sampled.

15.2. Comparison with Archetypal Analysis

Archetypal Analysis learns a global convex model for the data. This is appropriate when the data manifold is a convex subset of Euclidean space. In other situations the Archetypal Analysis fails to learn the details of the data. Figure 34 shows some structures that might be constructed by Archetypal Analysis.

In contrast, Activated Simplices learns a local convex model for the data. Figure 35 shows structures that might be learned by Activated Simplices.

Archetypal Analysis learns convex extremes but it can fail to find intrinsic boundaries. Activated Simplices can find convex extremes among the bases (though not all bases will be extremes), and it can find intrinsic boundaries when the manifold is well reconstructed.

References

- [1] Semeion handwritten digit data set. <http://archive.ics.uci.edu/ml/datasets/semeion+handwritten+digit>. 6
- [2] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *SCG '00: Proceedings of the sixteenth annual symposium on Computational geometry*, pages 1–10. ACM Request Permissions, May 2000. 14




Figure 36: Simplicial structures learned from data sampled from the torus, using $p = 4, 40$ and 60 bases and using $r = 0.1, 0.4, 0.8$ and 1 . $r = 0.1$ is heavily penalized and $r = 1$ is unpenalized.

- [3] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003. [1](#)
- [4] E. Bradley, H. Trevor, and T. Roberts. Least angle regression. *The Annals of Statistics*, 32(2):407–499, 2004. [4](#)
- [5] P. S. Bradley and O. L. Mangasarian. k-plane clustering. *Journal of Global Optimization.*, 16(1):23–32, 2000. [1](#)
- [6] K. Q. Brown. Voronoi diagrams from convex hulls. *Information Processing Letters*, 9(5):223–228, Jan. 1979. [12](#)
- [7] G. D. Canas, T. A. Poggio, and L. Rosasco. Learning Manifolds with K-Means and K-Flats. *NIPS*, 2012. [1](#), [2](#)
- [8] Y. Chen, J. Mairal, and Z. Harchaoui. Fast and robust archetypal analysis for representation learning. *arXiv preprint arXiv:1405.6472*, 2014. [2](#), [6](#), [15](#)




Figure 32: The red line is an interior simplex, the blue segments are boundary simplices



Figure 33: Each show shows a boundary simplex found by our method. We can see that, for most case, there is at least an “extreme” lighting condition in the simplex.




Figure 34: Bases that might be learned with Archetypal Analysis. The black line represents the data manifold and the red dots are the learned bases. (a) Archetype analysis will obtain similar bases as our approach when data is on boundary of convex set. (b,c,d) Archetype Analysis doesn't model details away from convex extremes. Note also that the regions that are represented by the bases (the convex hull of the bases) are much larger than the manifolds from which they are learned.

[9] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold recon-




Figure 35: Bases that might be learned by activated simplices. The black line represents the data manifold and the red dots are the learned bases. The red segments are the activated simplices.

struction from point samples. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, Jan. 2005. 14

[10] A. Cutler and L. Breiman. Archetypal analysis. *Technometrics*, 36(4):338–347, 1994. 2, 8, 15

[11] T. K. Dey and S. Goswami. Tight cocone: a water-tight surface reconstructor. In *SM '03: Proceedings of the eighth ACM symposium on Solid modeling and applications*. ACM Request Permissions, June 2003. 14

[12] D. L. Donoho. Neighborly polytopes and sparse solutions of underdetermined linear equations. Technical report, 2005. 10

[13] D. L. Donoho and C. Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences of the United States of America*, 100(10):5591–5596, 2003. 1

[14] H. Edelsbrunner. Shape reconstruction with Delaunay complex. pages 119–132, 1998. 14

[15] H. Edelsbrunner. Surface reconstruction by wrapping finite sets in space. pages 379–404, 2003. 14

[16] H. Edelsbrunner, D. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *Information Theory, IEEE Transactions on*, 29(4):551–559, 1983. 14

[17] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete & Computational Geometry*, 1(1):25–44, 1986. 13, 14

[18] E. Elhamifar and R. Vidal. Sparse subspace clustering: Algorithm, theory, and applications. *ieeexplore.ieee.org*, 2013. 1, 2

[19] J. Ham, D. D. Lee, S. Mika, and B. Schölkopf. A kernel view of the dimensionality reduction of manifolds. page 47, 2004. 1

[20] P. S. Huggins. *Sparse Coding via Geometry*. PhD thesis, Yale University, May 2005. 4, 10

[21] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, jul 2014. 6, 7

[22] W. Li, Z. Zhang, and Z. Liu. Action recognition based on a bag of 3d points. In *(CVPRW), 2010 IEEE Computer Society Conference on*, pages 9–14. IEEE, 2010. 6, 7

[23] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online dictionary learning for sparse coding. pages 689–696, 2009. 5, 8

- [24] T. Minka. Estimating a dirichlet distribution, 2000. [8](#)
- [25] B. A. Olshausen and D. J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996. [10](#)
- [26] M. R. Osborne, B. Presnell, and B. A. Turlach. On the LASSO and its dual. *Journal of Computational and Graphical Statistics*, 9(2):319–337, 2000. [4](#), [10](#)
- [27] N. Pitelis, C. Russell, and L. Agapito. Learning a Manifold as an Atlas. In *CVPR*, pages 1642–1649, 2013. [1](#), [2](#), [6](#)
- [28] N. Pitelis, C. Russell, and L. Agapito. Learning a Manifold as an Atlas Supplementary Material. *CVPR*, 2013. [6](#)
- [29] L. K. Saul and S. T. Roweis. Think globally, fit locally: unsupervised learning of low dimensional manifolds. *The Journal of Machine Learning Research*, 2003. [1](#)
- [30] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. [1](#)
- [31] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B. Methodological*, pages 267–288, 1996. [4](#), [10](#)
- [32] C. Wang, Y. Wang, Z. Lin, A. L. Yuille, and W. Gao. Robust Estimation of 3D Human Poses from a Single Image. In *CVPR*, 2014. [4](#), [7](#)
- [33] C. Wang, Y. Wang, and A. L. Yuille. An Approach to Pose-Based Action Recognition. In *CVPR*, pages 915–922, 2013. [7](#)
- [34] J. Wang, Z. Liu, Y. Wu, and J. Yuan. Mining actionlet ensemble for action recognition with depth cameras. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 1290–1297. IEEE, 2012. [7](#)
- [35] S. Wright and J. Nocedal. *Numerical optimization*, volume 2. Springer New York, 1999. [5](#)
- [36] Z. Zhang and H. Zha. Principal Manifolds and Nonlinear Dimensionality Reduction via Tangent Space Alignment. *SIAM journal on scientific computing*, 26(1):313–338, Jan. 2004. [6](#)